

# *Fast-Tracking* AI Value Creation in Financial Institutions

How prebuilt AI solutions mitigate risks, free up resources, and accelerate ROI while laying the foundation for enterprise-wide technology adoption.

## Key Takeaways

- ✓ Financial institutions looking to adopt generative AI must decide between building in-house solutions or purchasing third-party technology. While building may seem straightforward, most in-house projects fail to scale to production.
- ✓ Building GenAI in-house poses risks from technical complexity, scalability challenges, and unforeseen costs. Partnering with an experienced supplier accelerates development, frees IT teams, and reduces uncertainty.
- ✓ The Squirro Enterprise GenAI platform offers a robust foundation for scalable enterprise AI, enabling faster deployment and delivering measurable value without the delays and inefficiencies of building from scratch.



# The Mirage of Easy Enterprise AI Development

Does this sound familiar? Your organization is excited by the transformative potential of AI – maybe you’ve seen competitors use it to drive their efficiency, or you’ve experienced its potential in other areas of life. Either way, you recognize AI as the enabling technology your business needs to stay ahead or even leapfrog the competition.

Fueled by this vision, your IT team pushes forward and launches a proof of concept (PoC) to test the waters. The initial results come in: The technology shows promise, and everyone is eager to move forward.

But then reality sets in. Technical hurdles, integration issues, and scaling challenges drag out deployment, consuming more time and resources than expected while operational risks related to privacy, security, and regulatory compliance become apparent. Despite the early success, the solution never makes it to production, with costly consequences: lost time, wasted resources, exceeded budgets, and missed business opportunities.

Worse yet, lost confidence in AI: What was once an exciting path to competitive advantage now feels like a risky gamble. And as the clock ticks, you fall further behind competitors who successfully leverage AI to drive innovation and differentiation, leaving you wondering: How did the dream of AI become a missed opportunity?

# Defining Requirements: Technology, Scope, and Scale

The decision to build or buy generative AI technology depends on aligning your organization's needs, as well as its intended scope and scale, with available resources and strategic priorities. For financial institutions, this typically means choosing between fine-tuning large language models or retrieval augmented generation (RAG), evaluating single-use applications versus enterprise-wide platforms, and ensuring scalability to meet growing demands.

By carefully assessing these dimensions, organizations can chart a clear path to AI adoption that supports strategic goals and delivers sustainable value.

## Technology

**Fine-tuning** foundational large language models has been used to develop tailored solutions. Still, it demands significant resources and delivers lower accuracy and security, e.g., due to outdated training data and the risk of exposing sensitive information.

**Retrieval augmented generation**, or RAG, provides generative AI based on securely sourced, real-time enterprise data. RAG can be used in conjunction with access control lists (ACLs) while avoiding the high cost of retraining models.

**Enhanced RAG**, extends RAG's foundational capabilities with knowledge graphs, AI guardrails, an agent framework, data virtualization, and a privacy layer, further increasing accuracy, contextual understanding, data access, compliance, and security.

## Scope

Single-use AI applications deliver quick wins but limit broader impact. Organization-wide AI platforms integrate AI across functions, unlocking synergies like optimized risk management and compliance.

## Scale

Vertical scaling handles growth within a use case, while horizontal scaling expands AI across departments. Scalable platforms adapt seamlessly to evolving needs, ensuring sustained value.

# Quantifying Success: Daily Victories, Key Metrics, and ROI


Because of the horizontal opportunities offered by the technology, all financial institutions are likely to integrate enterprise GenAI into their activities in the coming years – across departments, use cases, and user groups. However, successful implementation isn't just about deploying technology—it's about solving real business challenges. This includes **saving time** by automating tasks, **gaining efficiency** through instant data retrieval, and **boosting productivity** by streamlining complex workflows.

Technically, success can be defined by various **key metrics**:

- ✓ **Accuracy:** Ensuring the AI solution delivers precise, reliable insights that drive confident decision-making.
- ✓ **% of Enterprise Data Integrated:** The ability to seamlessly integrate diverse data sources across the organization, unlocking a full spectrum of insights.
- ✓ **Time to Deployment:** The speed with which AI is implemented, allowing the organization to start seeing results quickly.
- ✓ **Cost of Deployment:** Ensuring that the investment required for AI adoption is justified by the return on investment (ROI) achieved.
- ✓ **Security:** The platform's ability to safeguard sensitive data and comply with industry standards, mitigating risks.
- ✓ **Maintenance Efforts:** The ongoing effort needed to keep the solution running smoothly, ensuring sustainability and long-term value.
- ✓ **Compliance:** The ability of the AI solution to meet regulatory requirements, especially in heavily regulated industries like finance.
- ✓ **Risk Mitigation:** How well the solution helps mitigate business risks, such as compliance violations or missed opportunities.
- ✓ **User Adoption:** The rate at which employees embrace and effectively use the AI solution, directly impacting its success across the organization.
- ✓ **Scalability:** Both in terms of handling larger volumes of data and extending the AI solution horizontally across the organization to create broader value.

Ultimately, business success boils down to two critical factors: the **time it takes to generate value** and the **amount of value generated**. The faster and more impactful the AI solution, the greater the **return on investment**.





## Ready to Build? A Self-Assessment for In-House GenAI Development

Whether to build or buy your enterprise AI is a consequential decision requiring careful consideration. While it's safe to say that most financial institutions underestimate the complexity of in-house generative AI development and the opportunity costs involved, some may have the resources and expertise required to build their own solutions from scratch.

To help you evaluate whether your organization is prepared for the challenges involved, we've prepared a short self-assessment that will guide you through the critical factors to consider before committing to either path.

## Expectations

- Scope of the Solution:** First, consider your long-term strategy. Are you looking to develop a single-use AI application for a specific use case or an enterprise AI platform for your entire organization?
- Time to Deployment:** How soon do you need your solution to be operational? In-house development timelines often stretch beyond expectations due to challenges that are unique to AI and RAG systems.

## Capabilities

- Internal AI Expertise:** Does your IT team have the expertise to build, maintain, and optimize complex AI models at scale, including tasks like vectorization, indexing, and continuous updates?
- Data Management Capabilities:** Do you have the IT infrastructure needed to handle large-scale data ingestion and processing? AI solutions require seamless data flow from multiple sources, and, again, scalability is key.
- Security & Compliance:** Is your organization equipped to ensure data security and compliance? GenAI solutions must integrate secure access controls and meet stringent regulatory standards, particularly in financial institutions.
- Testing and Monitoring:** Do you have robust testing frameworks to continuously monitor AI performance and make adjustments as needed to maintain quality?
- Scalability:** Does your IT team have the technical experience required to build an AI platform that can scale both vertically (in terms of data size) and horizontally (across use cases departments)?
- Resource Availability:** Do you have the necessary resources (budget, personnel, time) to develop and maintain an AI solution in-house without neglecting other important tasks?
- ROI & Opportunity Costs:** Have you evaluated the potential return on investment and weighed it against the opportunity costs of building vs. buying an AI solution?

Should your internal capabilities and resources align with these challenges, building a GenAI solution in-house may be feasible – as long as the privacy, security, and compliance risks that the financial sector is subject to are well managed.

However, if you face significant gaps in resources, expertise, or scalability, a proven, prebuilt enterprise GenAI platform will offer a faster, more secure, and cost-effective path to success while freeing up your IT teams from unnecessary development tasks.

# Bridging the Gap Between Executive Vision and In-House Capabilities

In our experience, discussions around buying prebuilt enterprise AI typically begin the moment internal development efforts hit their limits and vast resources, in terms of time, money, and talent, have been spent.

*We underestimated the challenges of scaling in-house GenAI solutions, which led to unexpected delays and escalating maintenance costs. Partnering with Squirro transformed our approach – minimizing risks, accelerating time to value, and ensuring predictable, manageable maintenance costs.*

Customer feedback from a financial services firm

*While building and deploying individual GenAI applications in-house seemed manageable at first, each new project required starting from scratch – draining time, effort, and resources. Squirro's scalable GenAI platform allowed us to seamlessly expand our capabilities without reinventing the wheel for every additional use case.*

Customer feedback from a European bank

Developing a single-use AI-powered tool that draws on a static data repository comprising thousands of static documents may be manageable. Building a scalable enterprise-grade AI platform capable of harnessing millions of live documents spread across multiple data repositories, with thousands of users and access control lists, can become a drain on internal resources, risking to derail other IT projects and everyday operations.



# Risks and Obstacles in In-House AI Development

## Strategic Risks

- ✓ **Total Cost of Ownership (TCOO):** The high costs associated with building and maintaining AI systems may exceed expectations, affecting the overall budget.
- ✓ **Time to Deployment:** Delays in rolling out the AI solution prevent businesses from gaining timely insights and competitive advantages.
- ✓ **ROI Targets:** Failure to meet expected return on investment, as the platform may take longer to deliver value or fall short in its capabilities.
- ✓ **Regulatory Compliance:** Failing to enforce strict regulatory compliance in the financial sector poses a strategic risk, jeopardizing long-term operational success.

## Technical Risks

- ✓ **Data Ingestion:** Failure to meet performance targets due to challenges in efficiently processing and integrating vast amounts of data.
- ✓ **Information Retrieval Stack:** Inefficient searching, indexing, and data retrieval across systems result in slow responses, inaccuracies, and inefficiencies.
- ✓ **Scalability & Maintenance:** Struggles with scaling AI systems in size and across the organization, impacting accuracy and driving up operational costs (OPEX).
- ✓ **Security & Compliance:** Mismanagement of security protocols, access control lists (ACLs), and regulatory requirements, compromising data security and compliance.

## Operational Risks

- ✓ **Accuracy and Performance Gaps:** In-house solutions struggle to deliver consistent accuracy, impacting decision-making and overall performance.
- ✓ **Customization Challenges:** While initial tools may address specific tasks, meeting evolving enterprise-wide needs often proves resource-intensive and unsustainable.
- ✓ **Maintenance Overload:** Ongoing maintenance, scaling, and updates create a heavy operational burden, driving up OPEX and diverting resources from strategic priorities.
- ✓ **Resource Constraints:** Building in-house strains development teams, requiring talent, time, and budgets that often exceed initial expectations and delay time-to-value.

# Mapping Out the *Benefits* of Buying Early

For business leaders, solution architects, technological innovators, and operational experts in the financial sector, the decision to buy an enterprise GenAI platform early rather than attempting to build one in-house from scratch offers a variety of advantages.

By grounding enterprise AI applications on a prebuilt platform, **strategic leaders** can fast-track their AI initiatives, ensuring they stay competitive and achieve quicker ROI.

**Technical teams** can avoid the burden and complexity of building and maintaining AI from scratch, allowing them to focus on optimizing and innovating other key areas.

**Innovation leaders** gain the ability to deliver impactful AI solutions rapidly, improving user experiences without the lengthy development cycle.

Meanwhile, **operational experts** benefit from smoother, more predictable deployments, reducing risk and ensuring scalability across the organization.

Ultimately, leveraging a prebuilt platform as a baseline accelerator early on in their AI development journey enables financial institutions to accelerate time-to-value, mitigate operational disruptions, and deliver tangible business results faster.



# *Fast-Tracking AI Value Creation with Squirro*

## **Deep Industry Expertise and Experience**

Squirro brings years of deep industry knowledge and a proven track record in delivering AI-driven solutions tailored to your specific needs. We understand the unique challenges of heavily regulated industries, particularly in sectors like finance, and can build solutions that integrate seamlessly with existing systems while addressing your critical business requirements.

## **Beyond Basic RAG: Enhanced AI Capabilities**

Squirro's AI solutions go beyond traditional retrieval-augmented generation (RAG) models. With enhanced RAG, Squirro optimizes how enterprise AI retrieves and generates information, improving the accuracy and relevance of insights. This helps your organization unlock deeper insights faster, driving smarter decision-making with greater confidence.

## **Cost-Effectiveness at Scale**

As your organization grows, our solutions scale cost-effectively without compromising on performance. By reducing the overhead and complexity of maintaining in-house AI solutions, Squirro allows you to focus resources on strategic growth initiatives while keeping costs in check.

## **Effective Access Control List (ACL) Management**

Squirro offers robust access control list (ACL) management at scale to ensure that sensitive data is handled securely and in compliance with industry standards. This maintains a secure, regulated environment while enabling appropriate access for different teams, fostering trust and reducing security risks.

## **A Platform That Meets Your Organization's Full Needs**

Squirro helps you build a comprehensive platform that addresses the full spectrum of your organization's needs. From data management to analytics, the platform is designed to work across departments, providing a unified solution that enhances collaboration, streamlines workflows, and drives efficiency at every level of the organization.



## Summary

It should be clear by now that the question is not whether your organization will integrate GenAI into its operations but rather how soon you will benefit from its transformational benefits. Succeeding on this journey requires balancing the risk and opportunities of in-house and prebuilt solutions.

As laid out here, we've observed that most organizations fail to scale in-house GenAI initiatives to production because they overestimate their internal development capabilities and underestimate the long-term challenges involved. Ultimately, issues like data ingestion, scalability, and security often lead to missed performance targets and unmet ROI.

As Dorian Selz, co-founder and CEO of Squirroi puts it: "There is no replacement for experience in operating a vectorized index at scale with full permissions over time. It involves a steep learning curve, and those who try to do it themselves will pay a price."

By partnering with Squirro, you don't just get a platform that supports current needs – you unlock opportunities you may not have even considered. Our expertise and platform capabilities help you discover new use cases, innovate faster, and identify overlooked growth areas, keeping you ahead of the competition and maximizing your AI investments.

To learn more about how to leverage our prebuilt GenAI platform to mitigate risks, free up resources, and accelerate ROI in financial institutions:

[Get in Touch](#)

[Book a Demo](#)